

Fourier analysis of the CGMN method for solving the Helmholtz equation

Tristan van Leeuwen,

Dept. of Earth and Ocean sciences University of British Columbia Vancouver, BC, Canada

October 10, 2012

Abstract

The Helmholtz equation arises in many applications, such as seismic and medical imaging. These application are characterized by the need to propagate many wavelengths through an inhomogeneous medium. The typical size of the problems in 3D applications precludes the use of direct factorization to solve the equation and hence iterative methods are used in practice. For higher wavenumbers, the system becomes increasingly indefinite and thus good preconditioners need to be constructed. In this note we consider an accelerated Kazcmarz method (CGMN) and present an expression for the resulting iteration matrix. This iteration matrix can be used to analyze the convergence of the CGMN method. In particular, we present a Fourier analysis for the method applied to the 1D Helmholtz equation. This analysis suggests an optimal choice of the relaxation parameter. Finally, we present some numerical experiments.

1 Introduction

The Helmholtz equation arises in many applications such as seismic and medical imaging. These application are characterized by the need to propagate many wavelengths through an inhomogeneous medium. The scale of the problems in 3D application precludes the use of direct factorization to solve the equation and hence iterative methods are used in practice. For higher wavenumbers, the system becomes increasingly indefinite and thus good preconditioners need to be constructed. For an overview of preconditioning techniques for the Helmholtz equation we refer to [3] and references cited therein.

In this note, we consider an accelerated Kazcmarz method (CGMN, [1]) for solving the Helmholtz equation [4]. We present an expression of the iteration matrix that allows us to analyze its convergence behaviour.

We proceeds as follows. First, we give a brief overview of the CGMN method and its relation to SSOR. This gives us an expression for the iteration matrix in terms of the original matrix. Then, we present the Fourier analysis for a 1D finite-difference discretization of the Helmholtz equation and we present some numerical examples. Finally, we draw conclusions and discuss future work.

1.1 The CGMN method

The CGMN method [1] relies on a symmetric Kaczmarz (SKACZ) sweep in conjunction with the conjugate gradient (CG) method.

The Kaczmarz method solves a system of N equations, $A\mathbf{u} = \mathbf{s}$, by cyclically projecting the iterate onto rows of the matrix [5]

$$\mathbf{u} := \mathbf{u} + \omega (s_i - \mathbf{a}_i^T \mathbf{u}) \mathbf{a}_i / \|\mathbf{a}_i\|_2^2, \quad i = 1 \dots N, \quad (1)$$

where \mathbf{a}_i denotes the i -th row of A as column vector and $0 < \omega < 2$ is a relaxation parameter. Note that we may also let ω vary per row. Introducing the matrices $Q_i = (I - \omega \mathbf{a}_i \mathbf{a}_i^T / \|\mathbf{a}_i\|_2^2)$, we may write this iteration as

$$\mathbf{u} := Q_i \mathbf{u} + \omega \mathbf{a}_i s_i / \|\mathbf{a}_i\|_2^2 \quad (2)$$

A *double sweep* through the matrix (from row 1 to N and back) can then be denoted by

$$\mathbf{u} := Q\mathbf{u} + R\mathbf{s}, \quad (3)$$

where $Q = Q_0 Q_1 \dots Q_{N-1} Q_{N-1} \dots Q_0$ and R contains all the factors multiplying \mathbf{s} . It is easily verified that the matrices Q_i are symmetric and have rank 1 with eigenvalue $1 - \omega$. It follows that Q is symmetric and has eigenvalues $\in [-1, 1]$. Hence, $(I - Q)$ is symmetric and positive semi-definite and we can use CG to solve the equivalent system

$$(I - Q)\mathbf{u} = R\mathbf{s}. \quad (4)$$

The question is, how does this new system behave in terms of ω ? In the next section, we will first derive a convenient expression for the matrix Q in terms of the matrix A , and in the subsequent section we present a Fourier analysis of the error propagation applied to the 1D Helmholtz equation.

1.2 Relation to SSOR

It is well-known that a SKACZ sweep on the system $A\mathbf{u} = \mathbf{s}$ is equivalent to a SSOR sweep on the normal equations $AA^T \mathbf{x} = \mathbf{s}$, $\mathbf{u} = A^T \mathbf{x}$ [1]. Splitting the matrix AA^T into its diagonal, strictly lower and upper triangular parts D, L and L^T such that $AA^T = D + L + L^T$, the SSOR iteration can be represented as follows [7]

$$\mathbf{x} := G\mathbf{x} + H\mathbf{s}, \quad (5)$$

where

$$G = (D + \omega L^T)^{-1}((1 - \omega)D - \omega L)(D + \omega L)^{-1}((1 - \omega)D - \omega L^T), \quad (6)$$

$$H = \omega(2 - \omega)(D + \omega L^T)^{-1}D(D + \omega L)^{-1}. \quad (7)$$

Since the SKACZ and SSOR iterations are equivalent, the matrices Q, R and G, H are related via

$$QA^T = A^T G, \quad (8)$$

$$R = A^T H. \quad (9)$$

Next, we note that we may write G in terms of H and A as

$$G = I - HAA^T, \quad (10)$$

from which it follows that we can write

$$Q = I - A^T H A. \quad (11)$$

2 Fourier analysis of the error propagation

We study the properties of the preconditioner by looking at the *error propagation matrix*. This matrix relates errors in \mathbf{u} when doing a standard Richardson iteration on the preconditioned system:

$$\mathbf{u}_{k+1} = (I - (I - Q))\mathbf{u}_{k+1} + R\mathbf{b}. \quad (12)$$

In terms of the error $\mathbf{e}_{k+1} = \mathbf{u}_k - \mathbf{u}^*$, we get

$$\mathbf{e}_{k+1} = Q\mathbf{e}_k, \quad (13)$$

Similarly, the residual $\mathbf{r}_k = A\mathbf{e}_k$ is propagated by

$$\mathbf{r}_{k+1} = G^T \mathbf{r}_k. \quad (14)$$

We analyze the behaviour of the SKACZ preconditioner through a Fourier analysis of the error propagation matrix [2, 6]. We decompose the error into its Fourier modes $e_j(\theta) = e^{ij\theta}$ and want to find the amplitude $a(\theta)$ such that

$$Q\mathbf{e}(\theta) = a(\theta)\mathbf{e}(\theta), \quad (15)$$

Using eq. (7) and (11) we can factorize the amplitude function in terms of those of the matrices A (a_1), $(D + \omega L)$ (a_2), D (a_3) and $(D + \omega L^T)$ (a_4):

$$a = 1 - \omega(2 - \omega) \frac{a_1^2 a_3}{a_2 a_4}. \quad (16)$$

Note that the matrix G has the same amplitude function.

2.1 1D case

We consider a simple finite difference discretization of the the 1D Helmholtz equation $(k^2 + d^2/dx^2)u = s$ on the domain $\Omega = [0, 1]$ with Dirichlet boundary conditions. The resulting system of equations is denoted by

$$A\mathbf{u} = \mathbf{s}, \quad (17)$$

where

$$a_{ii} = k^2 - 2h^{-2}, \quad (18)$$

$$a_{i,i\pm 1} = h^{-2}, \quad (19)$$

and $h = 1/(N+1)$ denotes the gridspacing and N denotes the number of interior gridpoints. The matrix A is symmetric and has eigenvectors

$$\mathbf{v}_i^n = \sin(n\pi ih), \quad n = 1, 2, \dots, N \quad (20)$$

with corresponding eigenvalues

$$\lambda_n = k^2 + 2h^{-2}(\cos(n\pi h) - 1). \quad (21)$$

The entries of the matrices D and L are given by $d_{ii} = \|\mathbf{a}_i\|_2^2 = \gamma^2 + 2h^{-4}$, $l_{i,i-1} = \mathbf{a}_i^T \mathbf{a}_{i-1} = 2h^{-2}\gamma$ and $l_{i,i-2} = h^{-4}$ where $\gamma = k^2 - 2h^{-2}$.

It is now readily verified that

$$a_1 = \gamma + 2h^{-2} \cos \theta, \quad (22)$$

$$a_2 = \gamma^2 + 2h^{-4} + \omega h^{-2} (2\gamma \exp^{-i\theta} + h^{-2} \exp^{-2i\theta}), \quad (23)$$

$$a_3 = \gamma^2 + 2h^{-4}, \quad (24)$$

$$a_4 = \gamma^2 + 2h^{-4} + \omega h^{-2} (2\gamma \exp^{+i\theta} + h^{-2} \exp^{+2i\theta}). \quad (25)$$

Substituting these expressions in eq. (16) yields

$$a(\theta, \omega, h) = 1 - \frac{\omega(2 - \omega)\beta(\gamma + 2h^{-2} \cos \theta)^2}{\beta^2 + 2\beta\omega(2\gamma h^{-2} \cos \theta + h^{-4} \cos 2\theta) + \omega^2 h^{-4}(4\gamma^2 + 4\gamma h^{-2} \cos \theta + h^{-4})}, \quad (26)$$

where $\beta = \gamma^2 + 2h^{-4}$.

To ensure a certain amount of gridpoints per wavelength, n_g , we let $kh = 2\pi/n_g$. Note that in this case, the amplitude only depends on n_g and ω . The question is, what is the optimal ω for a given n_g . Figure 4 (a) shows the amplitude as a function of θ and ω for $n_g = 10$. In Figure 4 (b) we show $\frac{\max_{\theta}(1-a(\theta))}{\min_{\theta}(1-a(\theta))}$, which is roughly equivalent to plotting the condition number of $I - Q$, and this suggests an optimal value of $\omega \approx 1.5$ for $n_g = 10$. Figure 4 (c) shows the optimal ω as a function of n_g . We can use this curve to adapt ω to the (local) wavenumber.

3 Numerical experiments

We present some numerical experiments that verify the results discussed above.

3.1 1D

In the first experiment we solve the Helmholtz equation for various wavenumbers k , with a fixed number of $n_g = 10$ gridpoints per wavelength. Figure 4 (a) shows the number of CG iterations needed to converge to a tolerance of 10^{-6} for various ω . The predicted optimal value (cf. figure 4 (c)) is also indicated and nicely coincides with the optimal ω in terms of the iteration count.

In the second experiment, we use the same values of k , but choose the gridspacing based on the *highest* wavenumber used. This way, the number of gridpoints per wavelength is much higher for the lower wavenumbers. The result is shown in Figure 4 (b). Again, the predicted optimal value of ω is close to the empirical optimum.

3.2 2D

We use a medium that consist of a high-wavenumber anomaly embedded in a background with constant wavenumber k_0 , as depicted in figure 4 (a). The wavefield resulting from an incident planewave $\exp(ik_0x)$ (figure 4 (b)) is depicted in figure 4 (c). The convergence histories using either a constant $\omega = 1.5$ (chosen according to the highest wavenumber) or $\omega = 1.95$ (chosen according to the lowest wavenumber) or a spatially varying ω (corresponding to the local wavenumber) are shown in figure 4 (c). It appears that for a medium with a large contrast it may indeed be beneficial to vary the relaxation parameter spatially. Interestingly enough, the convergence for $\omega = 1.95$ is much faster in the beginning, suggesting perhaps that it may be beneficial to change the relaxation parameter per iteration.

4 Conclusion and future work

We have presented a Fourier analysis of the CGMN method applied to the 1D Helmholtz equation. The amplitude function of the error propagation matrix suggests that the optimal relation parameter depends on the number of gridpoints per wavelength $n_g = 2\pi/(kh)$. This observation was confirmed by counting the number of CGMN iterations needed to actually solve the system for various values of ω and n_g . The same optimal relation can be used to adapt the relaxation parameter to the local wavenumber in case of inhomogeneous media. A numerical example on a 2D medium with a very high contrast suggests that this might indeed improve convergence.

Future research is aimed at extending this analysis to 2D/3D Helmholtz equations and varying media, as well as the parallel extension of the CGMN method presented by [4]. An optimal strategy to adapt the relaxation parameter to the medium and perhaps vary it per iteration is of particular interest.

Acknowledgments

The author thanks Dan and Rachel Gordon for valuable discussions on this topic. This work was in part financially supported by the Natural Sciences and Engineering Research Council of Canada Discovery Grant (22R81254) and the Collaborative Research and Development Grant DNOISE II (375142-08). This research was carried with support from the sponsors of the SINBAD consortium.

References

- [1] Åke Björck and Tommy Elfving. Accelerated projection methods for computing pseudoinverse solutions of systems of linear equations. *BIT*, 19(2):145–163, June 1979.
- [2] A. Brandt. Multi-level adaptive solutions to boundary-value problems. *Mathematics of Computation*, 31:333–390, 1977.
- [3] O. G. Ernst and M. J. Gander. *Why it is Difficult to Solve Helmholtz Problems with Classical Iterative Methods*, volume 83 of *Numerical Analysis of Multiscale Problems*. 2011.
- [4] Dan Gordon and Rachel Gordon. Robust and highly scalable parallel solution of the Helmholtz equation with large wave numbers. *Journal of Computational and Applied Mathematics*, 237(1):182–196, January 2013.
- [5] S. Kaczmarz. Angenäherte Auflösung von Systemen linearer Gleichungen. *Bulletin International de l’Académie Polonaise des Sciences et des Lettres*, 35:355–357, 1937.
- [6] Rob Kettler. Analysis and comparison of relaxation schemes in robust multigrid and preconditioned conjugate gradient methods. In W. Hackbusch and U. Trottenberg, editors, *Multigrid Methods*, volume 960 of *Lecture Notes in Mathematics*, pages 502–534. Springer Berlin / Heidelberg, 1982. 10.1007/BFb0069941.
- [7] Yousef Saad. *Iterative methods for sparse linear systems*. SIAM, 2 edition, 2003.




Figure 1: (a) Amplitude (cf. eq. 26) as a function of θ and ω for a fixed number of gridpoints per wavelength $n_g = 10$. (b) Shows the ‘condition number’ as a function of ω for $n_g = 10$. The optimal $\omega \approx 1.5$. (c) shows the optimal ω as a function of n_g .




Figure 2: Number of CG iterations needed to converge to a tolerance of 10^{-6} for various ω with either a fixed number of gridpoints per wavelength $n_g = 10$ (a) or a fixed gridspacing based on the highest wavenumber used (b). The predicted optimal ω is indicated by a dashed line and coincides with the lowest iteration count.

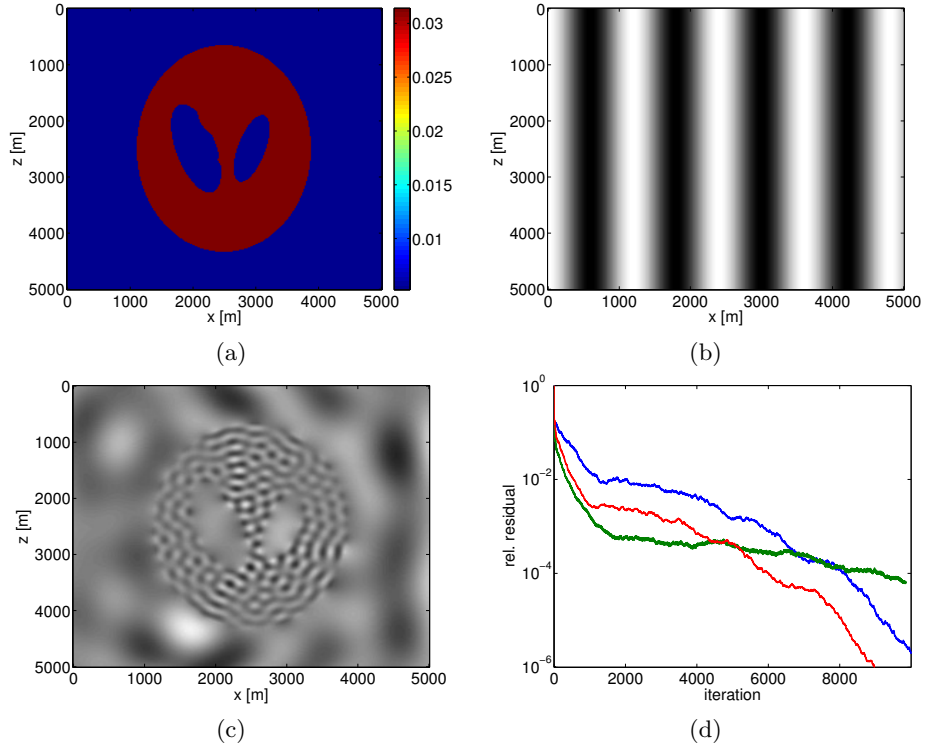


Figure 3: (a) wavenumber profile, (b) incident plane wave, (c) resulting wavefield and (d) convergence histories for $\omega = 1.5$ (blue), $\omega = 1.95$ (green) and a spatially varying ω (red).